

AD-A111 723

GEORGIA INST OF TECH ATLANTA SCHOOL OF INFORMATION A--ETC F/6 9/2
RESEARCH PROGRAM IN FULLY DISTRIBUTED PROCESSING SYSTEMS.(U)
OCT 81

N00014-79-C-0873

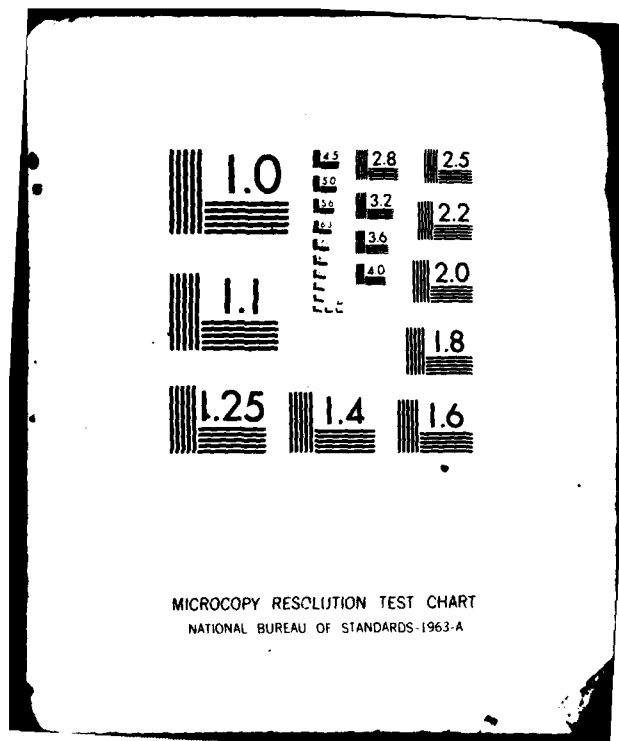
NL

UNCLASSIFIED

1-1
1-1



END
DATE
FILMED
4-82
NTIC



12

ADA111723



DTIC
ELECTE
MAR 9 1982

School of

Information and Computer Science

GEORGIA INSTITUTE OF TECHNOLOGY

A UNIT OF THE UNIVERSITY SYSTEM OF GEORGIA

DTIC FILE COPY

Approved for public release
Distribution Unlimited

THE GEORGIA INSTITUTE OF TECHNOLOGY
RESEARCH PROGRAM IN
FULLY DISTRIBUTED PROCESSING SYSTEMS

Quarterly Progress Report Number 8
1 June, 1981 - 31 August, 1981

October, 1981

Supported by

Office of Naval Research (ONR)
Contract: N00014-79-C-0873
GIT Project: G36-643

U.S. Air Force Rome Air Development Center (RADC)
Contract: F30602-78-C-0120
GIT Project: G36-654

U.S. Air Force Rome Air Development Center (RADC)
Contract: F30602-81-C-0249
GIT Project: G36-659

U.S. Army Research Office (ARO)
Contract: DAAG29-79-C-0155
GIT Project: G36-638

U.S. Army Institute for Research in
Management Information and Computer Science (AIRMICS)
Contract: DAAK70-79-D-0087
GIT Project: G36-647

National Science Foundation (NSF)
Contract: MCS-7924370
GIT Project: G36-652

School of Information and Computer Science
Georgia Institute of Technology
Atlanta, Georgia 30332

1. INTRODUCTION

This is the Eighth Quarterly Progress Report prepared covering the Georgia Tech Research Program in Fully Distributed Processing Systems (FDPS).

a. Program Description.

The Georgia Tech Research Program in Fully Distributed Processing Systems is a comprehensive investigation of data processing systems in which both the physical and logical components are extremely loosely coupled while operating with a high degree of control autonomy at the component level. The definition of the specific class of multiple computer systems being investigated, and the operational characteristics and features of those systems is motivated by the desire to advance the state-of-the-art for that class of systems that will deliver a high proportion of the benefits currently being claimed for distributed processing systems. The scope of individual topics being investigated under this program ranges from formal modeling and theoretical studies to empirical examinations of prototype systems and simulation models. Also included within the scope of the program are areas such as the utilization of FDPS's and their interaction with management operations and structure.

b. Program Support.

The principle support for the program is a Selected Research Opportunity contract from the Office of Naval Research; however, there are a number of other sources of funding which also support the program. A list of the currently active contracts and grants is given below.

Title: "Research on Fully Distributed Data Processing Systems"
Funding Agency: Office of Naval Research (ONR)
Contract Number: N00014-79-C-0873
GIT Project No.: G36-643/336
Principle Investigator: Philip H. Enslow, Jr.

Title: "Evaluation of Distributed Control Models"
Funding Agency: U.S. Air Force Rome Air Development Center (RADC)
Contract Number: F30602-78-C-0120
GIT Project No.: G36-654
Principle Investigator: Philip H. Enslow, Jr.

Title: "System Support Capabilities for Fully-Distributed Loosely-Coupled Processing Systems"
Funding Agency: U.S. Air Force Rome Air Development Center (RADC)
Contract Number: F30602-81-C-0249
GIT Project No.: G36-659
Principle Investigator: Philip H. Enslow, Jr.

Title: "Theory of Systems of Asynchronous Parallel Processors"
Funding Agency: U.S. Army Research Office (ARO)
Contract Number: DAA029-79-C-0155
GIT Project Number: G36-638/332
Principle Investigator: Nancy A. Lynch

Accession
DTIC TAB
Unannounced
Justification
By
Distribution/
Availability Codes
Dist A and/or
Special

1000
1000
1000

Title: "Support of MILPERCEN Data Storage Concept"
Funding Agency: U.S. Army Institute for Research in Management Information
and Computer Science (AIRMICS)
Contract Number: DAAK70-79-D-0087
GIT Project Number: G36-647
Principle Investigator: Alton P. Jensen

Title: "Complexity and Computability for Distributed Data Bases"
Funding Agency: National Science Foundation (NSF)
Contract Number: MCS-7924370
GIT Project Number: G36-652/340
Principle Investigator: Nancy A. Lynch

2. ORGANIZATION AND STAFFING

Faculty

DeMillo, Richard A. - Professor
Enslow, Philip H. Jr. - Professor
Griffeth, Nancy - Assistant Professor
Jensen, Alton P. - Professor
LeBlanc, Richard - Assistant Professor
Livesey, Jon - Assistant Professor
Lynch, Nancy A. - Associate Professor
Miller, Raymond - Professor
Underwood, William - Assistant Professor

Staff

McDonell, Sharon - Senior Secretary
Myers, Jeanette - Research Scientist
Pinion, Nancy - Part-time Secretary
Mongiovi, Roy - Research Technologist I

Students

There are approximately 30 students working on various projects in the FDPS Research Program. Of these, 12 are in the Ph.D. program, and 5 are preparing their M.S. Thesis on topics in FDPS.

3. CURRENT RESEARCH PROJECTS

The specific research projects have been organized into the major areas identified in the basic program proposal.

A. Theoretical and Formal Studies

- A.2 Decomposition of Parallel Systems
- A.3 Reliable Systems
- A.4 Time Performance of Distributed Systems
- A.5 Audit Algorithms
- A.6 Ticket Systems
- A.7 Synchronous Simulation
- A.8 Distributed Resource Allocation
- A.9 Theory of Distributed Databases
- A.10 Arbiter Design
- A.14 Using Complementary Distributed System Models
- A.15 Probabilistic Algorithms in Distributed Systems
- A.16 Stochastic Synchronization
- A.17 Research Allocation in a Failure-Prone Environment
- A.18 Multilevel Atomicity

B. Physical Interconnection and Networking

- B.2 Local Networking in Fully Distributed Processing Systems

C. Distributed Operating Systems

- C.1 Decentralized and Distributed Control
- C.2 Resource Allocation and Work Distribution in an FDPS
- C.4 Local Operating System
- C.5 Communications Support for Distributed Systems
- C.8 Distributed Software Tools
- C.9 Command Languages in an FDPS

D. Distributed Data Bases

- D.1 Concurrency Control in Distributed Database Systems
- D.2 Support of MILPERCEN Data Storage Concept
- D.3 Implementation of the Audit Algorithm

E. Fault-Tolerance

F. Special Hardware to Support FDPS

G. Application of Distributed Processing

H. System Design Methodologies

H.2 Coordinating Large Programming Projects

I. System Utilization

I.1 A Language for Distributed Programming
I.2 System Implementation Language Development
I.3 Experiments with a Distributed Compiler

J. Security

J.1 Process Structures
J.2 System Security

K. System Management

L. Evaluation and Comparison

L.1 Simulation of Distributed Algorithms (Griffeth, Lynch)

M. FDPS Testbed

M.1 Establishment of FDPS Testbed Facility
M.2 Remote Load Emulator
M.3 Fully Distributed Operating System Simulation Testbed

4. SUMMARY OF PROGRESS

A.2 Decomposition of Parallel Systems (Lynch, Fischer)

No significant progress to report.

A.3 Reliable Systems (Lynch, Fischer, Lamport, Merritt)

A paper, "A Lower Bound on Time to Achieve Interactive Consistency", was written, describing lower bound results for number of rounds to solve the Byzantine Generals problem. Mike Merritt generalized the result to an environment allowing authenticated communication.

A.4 Time Performance of Distributed Systems (Lynch, Fischer, Lazowska, Schönhage)

No significant progress to report.

A.5 Audit Algorithms (Griffeth, Fischer, Lynch)

Conditions under which the audit can be significantly optimized were identified. These are: (1) transactions visit the same sites regardless of the interleaving of transactions, and (2) transactions access the same data-items regardless of the interleaving of transactions. Complexity analysis for the normal case and these two special cases was done.

A paper, "Global States of a Distributed System", was presented, describing the general checkpoint algorithm. The paper was invited for submission to the special issue of TOSE based on the conference, and has been submitted. It is also being recast in database terms for submission to TODS, including some new results on complexity analysis.

A.6 Ticket Systems (Fischer, Griffeth, Guibas, Lynch)

A new algorithm was devised for dynamic tree-balancing when ticket returns are allowed. The tickets are moved from one location to another in the tree using a heuristic which is similar to the heuristic for buyers. In the case of tickets, a ticket is sent out in some direction if that direction is "deficient" in tickets according to a local estimate.

The ticket system simulation was modified to use a regenerative approach to the choice of epochs. This eliminates dependencies of results on behavior preceeding or following the epoch. It also simplifies the statistical inference methods required to assess the accuracy.

New results have been obtained generalizing the previous analysis of sequential resource allocation and general (interfering) resource allocation. Results are currently being written up in a paper, "Expected Time Analysis of a Distributed Resource-Allocation Algorithm".

A.7 Synchronous Simulation (Lynch, Fischer, Arjomandi)

A paper, "A Difference in Efficiency between Synchronous and Asynchronous Systems", has been rewritten from the conference version for journal submission.

A.8 Distributed Resource Allocation (Lynch)

This project has been completed.

A.9 Theory of Distributed Databases (Lynch, Griffeth)

Discussions were begun attempting to integrate ideas about multilevel atomicity of transactions, multilevel data structuring, and multilevel consistency constraints in databases. Some ideas for concurrency control designs were also discussed.

A.10 Arbiter Design (Lynch, Griffeth, Schönhage, Fischer)

No significant progress to report.

A.14 Using Complementary Distributed System Models (Lynch, Rounds, R. Miller)

No significant progress to report.

A.15 Probabilistic Algorithms in Distributed Systems (Lynch, Arjomandi, Fischer)

No further progress anticipated. Project completed.

A.16 Stochastic Synchronization (DeMillo, R. Miller, Lipton)

No significant progress to report.

A.17 Resource Allocation in a Failure-Prone Environment (Fischer, Lynch, Burns, Borodin)

No significant progress to report.

A.18 Multilevel Atomicity (Lynch)

The paper, "Multilevel Atomicity: A new Correctness Criterion for Distributed Databases", was written and submitted for journal publication. It contains a preliminary proposal for a systematic way in which one might weaken the usual "serializability" definition for transactions, in order to obtain increased concurrency.

B.2 Local Networking in FDPSs (Enslow)

A report on initial work in this area has been prepared.

C.1 Decentralized and Distributed Control (Enslow, LeBlanc, Saponas)

The second phase of this project was completed on 30 June 1981. A final report entitled, "Performance of Distributed and Decentralized Control Models for Fully Distributed Processing Systems - Initial Simulation Studies", has been completed and submitted to the sponsors for review. Final publication of this report is pending the reception of sponsor's approval. Meanwhile, further simulation experiments which were suggested by the previous work are being conducted.

C.2 Resource Allocation and Work Distribution in an FDPS (Enslow, Sharp)

Evaluation of models of work distribution and resource allocation will be done by means of simulation. Work this quarter consisted of coding three

RAWD algorithms in PASCAL for inclusion in the simulator used in Project C.1, thoroughly testing each of the three algorithms, and preparing data to be used as input to the simulation experiment.

C.4 Local Operating System (Livesey, LeBlanc, Spafford, Myers, Fukuoka, Pitts)

Further work has taken place in investigations into appropriate Local Operating Systems structures. A report has been written on the meta-system (PRIMOS) approach.

The members of this project have been working on Distributed Software Tools, Project C.8. Previous work done for this project, especially the exploration of methods of adapting the PRIMOS operating system to act as a LOS prototype, is providing design and implementation support for the DSWT project.

C.5 Communications Support for Distributed Systems (Enslow, Skowbo)

The working draft of a project proposal has been prepared and is being expanded as a result of continuing efforts to identify and describe communications requirements for fully distributed processing systems. The fundamental importance of broadcast services to support distributed application processes has prompted a comprehensive investigation of transmission hardware, network topologies, and the access and control procedures to support these services. For long-haul networks, satellites continue to dominate the picture. Several local-area network technologies are being studied and compared in light of their unique advantages and limitations for locally distributed processing.

C.8 Distributed Software Tools (Myers, Livesey, Hopkins, Lee, McGraw, Fox)

Work continues on the Distributed Software Tools Project initiated last quarter. The first phase of design and implementation of DSWT involves creating a mechanism for remote process execution. This capability will serve as a tool that can be used to write many network applications as small command language programs. For example, a distributed mail facility could implement local mail facilities on each node of the network and direct that mail to the user's terminal with the command, "mail; mail@gt.b; mail@gt.c", assuming that the user is connected to gt.a and that gt.b and gt.c are the only other nodes in the network where mail can be stored for the user. As of now, DSWT consists of five components which are in various stages of completion: the DSWT command interpreter, server processes called hosts which are responsible for initiating remote process execution, remote I/O processes which serve as interfaces between the local user and a remote process, a network file system, and a message passing facility.

C.9 Command Languages in an FDPS (Badre, Myers, Greene)

This project, initiated this quarter, is divided into two subprojects. The first has as its objective the design of a "friendly" command language suitable for a distributed environment. Work this quarter has been

devoted to the collection and review of the available literature on command languages.

The second subproject, scheduled to begin next quarter, will attempt to determine if multiple command languages can and should be made available in an FDPS. It will also explore various ways in which multiple command languages can be integrated in a network to provide some (perhaps all) of the claimed benefits of an FDPS.

D.1 Concurrency Control in Distributed Database Systems (Griffeth, Livesey, Lynch)

Experiments were designed for studying concurrency control algorithms on the simulation tool developed for Project L.1. For this purpose, four parameters were identified: (1) method of guaranteeing consistency, (2) method of conflict resolution, (3) locus of control, and (4) management of deadlocks. Ten significant experiments were identified.

D.2 Support of MILPERCEN Data Storage Concept (Jensen, Doyle, Gehl, Bingham)

The final report has been written and is being prepared for publication.

D.3 Implementation of the Audit Algorithm (Griffeth, Livesey, Lynch)

No significant progress to report.

H.2 Coordinating Large Programming Projects (Enslow, Underwood, Smith)

Major research activities this quarter included investigation of large software development as problem-solving and as design activity. Papers were written on both topics. A major paper detailing proposed future research was completed by the end of the quarter.

I.1 A Language for Distributed Programming (LeBlanc, Maccabe)

Design work has been completed. An implementation using a number of compiler development tools, including the code generator developed under Project I.2, is currently in progress. A paper was prepared for submission to the Symposium on Principles of Programming Languages.

I.2 System Implementation Language Development (LeBlanc, Akin)

Allen Akin has completed his M.S. thesis work on the development of a reusable code generator for our PRIME Computers. This tool will enable us to build compilers which generate high-quality code without developing a customized code generator for each one. The code generator is currently being tested by about 20 students in a compiler class and extensions will be planned based on their experiences.

I.3 Experiments with a Distributed Compiler (LeBlanc, J. Miller)

John Miller is currently working on refinements of the experiments previously conducted as part of this project.

J.1 Process Structures (DeMillo, Lipton, R. Miller, Merritt, Thomas)

Michael Merritt and Barbara Smith Thomas are currently studying the application of cryptographic techniques to supply utilities in a distributed system. Recent results have centered around providing secure communications.

J.2 System Security (Livesey, Davida, DeMillo)

Further studies are being carried out into an architecture to support secure multiprogramming.

L.1 Simulation of Distributed Algorithms (Griffeth, Lynch)

The design and specification of a simulation tool for distributed database systems was completed. The tool was designed to be usable for general distributed algorithms also. Details can be found in the technical report.

M.1 Establishment of FDPS Testbed Facility (Myers, Mongiovi, Pitts, Fox)

The Distributed Software Tools (DSWT) Project initiated under this project number last quarter is now Project C.8. Although DSWT will enhance the testbed facility by providing users with a distributed subsystem and an extended capability for running concurrent programs, it was decided that it belonged in the Distributed Operating Systems effort since it is essentially a Network Operating System implemented on top of several local operating systems.

M.2 Remote Load Emulator (Myers, Enslow, Forsyth)

This project has been completed this quarter. The emulator and a user's guide are now available. Near term plans for the emulator involve performance testing for DSWT, Project C.8, and the local network, Project B.2.

M.3 FDOS Simulation Testbed (LeBlanc, Saponas, Myers)

The simulator is complete and has been used to validate the control models developed in Project C.1.

5. TRAVEL RELATED TO THE FDPS PROGRAM

Date of Trip: July, 1981 - August, 1982

Individual(s) Traveling: N. Lynch

Itinerary: Cambridge, Massachusetts

Purpose: On leave at MIT.

Date of Trip: 20-23 July, 1981

Individual(s) Traveling: N. Griffeth, N. Lynch

Itinerary: Pittsburg, Pennsylvania

Purpose: Attend IEEE Conference on Reliability in Distributed Software and Databases, and present paper, "Global States of a Distributed System".

Date of Trip: 3-7 August, 1981

Individual(s) Traveling: R. LeBlanc

Itinerary: Santa Cruz, California

Purpose: Attend a course in Functional Programming at the Institute in Computer Science at the University of California at Santa Cruz.

Date of Trip: 26-28 August, 1981

Individual(s) Traveling: R. DeMillo, M. Merritt

Itinerary: Santa Barbara, California

Purpose: Attend CRYPTO '81 Symposium.

Date of Trip: 30 August - 4 September, 1981

Individual(s) Traveling: N. Griffeth

Itinerary: Boston, Massachusetts (MIT)

Contact: N. Lynch

Purpose: Develop distributed algorithms for ticket systems and discuss further work on performance studies of ticket systems.

6. VISITORS

(no visitors to report for this quarter)

7. PUBLICATIONS

Author(s): A. Akin

Title: V-mode Code Generator User's Guide

Type: internal document

Status: printed

Publ. Date: June, 1981

Author(s): M. Fischer, N. Griffeth, and N. Lynch

Title: Global States of a Distributed System

Type: conference paper

Status: Presented at 1981 IEEE Conference on Distributed Software and Databases. Invited and submitted for publication in special issue of TOSE based on this conference. Database version in preparation for submission to TODS.

Publ. Date: July 21, 1981

Author(s): N. Griffeth

Title: A Simulation Tool for Distributed Database Systems

Type: technical report

Status: published

GIT Number: GIT-ICS-81/15

Publ. Date: August, 1981

Author(s): A. Akin

Title: A Reusable Code Generator for PRIME 50-series Computers

Type: M.S. Thesis

Status: being printed

GIT Number: GIT-ICS-81/16

Publ. Date: August, 1981

Author(s): D. Forsyth

Title: A Remote Terminal Emulator for PRIME Computers

Type: technical report

Status: being printed

GIT Number: GIT-ICS-81/12

Publ. Date: August, 1981

Author(s): D. Forsyth

Title: User's Guide for the PRIME Remote Terminal Emulator

Type: internal document

Status: printed

Publ. Date: August, 1981

Author(s): A.B. Maccabe and R.J. LeBlanc

Title: Communication Features for Distributed Computing Environments

Type: conference paper

Status: submitted to the Symposium on Principles of Programming Languages

Author(s): N. Lynch

Title: Multilevel Atomicity: A New Correctness Criterion for Distributed Databases.

Type: journal paper

Status: submitted for publication

GIT FDPS Research Program

Quarterly Prog Report 8

